

Active Inference or Control as Inference? A Unifying View

Abraham Imohiosen^{*†}, Joe Watson^{*§}, and Jan Peters[§]

[†]RWTH Aachen University, Germany

[§]IAS, Technical University Darmstadt, Germany

abraham.imohiosen@rwth-aachen.de

{watson,peters}@ias.informatik.tu-darmstadt.de

Abstract. Active inference (AI) is a persuasive theoretical framework from computational neuroscience that seeks to describe action and perception as inference-based computation. However, this framework has yet to provide practical sensorimotor control algorithms that are competitive with alternative approaches. In this work, we frame active inference through the lens of control as inference (CaI), a body of work that presents trajectory optimization as inference. From the wider view of ‘probabilistic numerics’, CaI offers principled, numerically robust optimal control solvers that provide uncertainty quantification, and can scale to nonlinear problems with approximate inference. We show that AI may be framed as partially-observed CaI when the cost function is defined specifically in the observation states.

Keywords: Active Inference · Control · Approximate Inference.

1 Introduction

Active inference (AI) [2, 4, 5] is a probabilistic framework for sensorimotor behavior that enjoyed sustained interest from computational neuroscientists. However, its formulation has been criticized for its opacity and similarity to optimal control [7–9], but is seemingly difficult to translate into an equally effective algorithmic form. In this work, we offer a critical analysis of AI from the view of control as inference (CaI) [1, 11, 14, 21, 24, 28], the synthesis of optimal control and approximate inference. The goal is to appreciate the insights from the AI literature, but in a form with computational and theoretical clarity.

2 Background

Here we outline the foundational theory and assumptions in this work.

2.1 Problem Formulation

We specifically consider a known stochastic, continuous, discrete-time, partially-observed, nonlinear, dynamical system with state $\mathbf{x} \in \mathbb{R}^{d_x}$, observations $\mathbf{y} \in$

\mathbb{R}^{d_y} and control inputs $\mathbf{u} \in \mathbb{R}^{d_u}$, operating over a time horizon T . We define the states in upper case to denote the variables over the time horizon, i.e. $\mathbf{U} = \{\mathbf{u}_0, \dots, \mathbf{u}_{T-1}\}$. The joint distribution (generative model) $p(\mathbf{Y}, \mathbf{X}, \mathbf{U})$ over these variables factorizes into several interpretable distributions: The dynamics $p(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t)$, observation model $p(\mathbf{y}_t | \mathbf{x}_t, \mathbf{u}_t)$, and behavior policy $p(\mathbf{u}_t | \mathbf{x}_t)$.

2.2 Variational Inference for Latent Variable Models

Inference may be described by minimizing the distance between the ‘true’ data distribution $p(\cdot)$ and a parameterized family $q_\theta(\cdot)$ [17]. A popular approach is to minimize the Kullback-Liebler (KL) divergence, e.g. $\min \mathbb{D}_{\text{KL}}[q_\theta || p]$ w.r.t. θ . More complex inference tasks can be described by observations \mathbf{y} influenced by unseen latent variables \mathbf{x} . Given an observation \mathbf{y}^* , maximizing the likelihood involves integrating over the hidden states, and so is termed the marginal likelihood $p(\mathbf{y}^*) = \int p(\mathbf{y}=\mathbf{y}^*, \mathbf{x})d\mathbf{x}$. Unfortunately this marginalization is typically intractable in closed-form. A more useful objective may be obtained by applying a variational approximation of latent state $q_\theta(\mathbf{x} | \mathbf{y}^*) = q_\theta(\mathbf{x} | \mathbf{y}=\mathbf{y}^*)$ to the log marginal likelihood and obtaining a lower bound via Jensen’s inequality [17]

$$\log \int p(\mathbf{y}^*, \mathbf{x})d\mathbf{x} = \log \int p(\mathbf{y}^*, \mathbf{x}) \frac{q_\theta(\mathbf{x} | \mathbf{y}^*)}{q_\theta(\mathbf{x} | \mathbf{y}^*)} d\mathbf{x} = \log \mathbb{E}_{\mathbf{x} \sim q_\theta(\cdot | \mathbf{y}^*)} \left[\frac{p(\mathbf{y}^*, \mathbf{x})}{q_\theta(\mathbf{x} | \mathbf{y}^*)} \right], \quad (1)$$

$$\geq \mathbb{E}_{\mathbf{x} \sim q_\theta(\cdot | \mathbf{y}^*)} \left[\log \frac{p(\mathbf{y}^*, \mathbf{x})}{q_\theta(\mathbf{x} | \mathbf{y}^*)} \right] = -\mathbb{D}_{\text{KL}}[q_\theta(\mathbf{x} | \mathbf{y}^*) || p(\mathbf{x}, \mathbf{y}^*)], \quad (2)$$

$$= \mathbb{E}_{\mathbf{x} \sim q_\theta(\cdot | \mathbf{y}^*)} [\log p(\mathbf{y}^* | \mathbf{x})] - \mathbb{D}_{\text{KL}}[q_\theta(\mathbf{x} | \mathbf{y}^*) || p(\mathbf{x})], \quad (3)$$

where equations 2, 3 are variations of the ‘evidence lower bound objective’ (ELBO). The expectation maximization algorithm (EM) [17], can be understood via Equation 3 as iteratively estimating the latent states (minimizing the KL term via q) in the E step and maximizing the likelihood term in the M step.

3 Active Inference

Active Inference frames sensorimotor behaviour as the goal of equilibrium between its current and desired observations, which in practice can be expressed as the minimization of a distance between these two quantities. This distance is expressed using the KL divergence, resulting in a variational free energy objective as described in Section 2.2. Curiously, AI is motivated directly by the ELBO, whose negative is referred to in the AI literature as the ‘free energy’ $\mathcal{F}(\cdot)$. The minimization of this quantity, $\mathcal{F}(\mathbf{y}^*, \mathbf{x}, \mathbf{u}) = \mathbb{D}_{\text{KL}}[q_\theta(\mathbf{x}, \mathbf{u} | \mathbf{y}^*) || p(\mathbf{x}, \mathbf{u}, \mathbf{y}^*)]$, as a model of behavior (i.e. state estimation and control), has been coined the ‘free energy principle’.

3.1 Free Energy of the Future

Despite the ELBO not being temporally restricted, AI delineates a ‘future’ free energy. This free energy is used to describe the distance between future predicted

and desired observations, where \mathbf{u} is directly represented as a policy $\mathbf{u} = \pi(\mathbf{x})$, so $\mathcal{F}(\mathbf{y}_t^*, \mathbf{x}_t | \pi)$ over the future trajectory is minimized. In active inference, π is commonly restricted to discrete actions or an ensemble of fixed policies, so inferring $p(\pi)$ can be approximated through a softmax $\sigma(\cdot)$ applied to the expected ‘future’ free energies for each policy over $t = [\tau, \dots, T - 1]$, with temperature γ and prior $p(\pi)$

$$p(\pi | \mathbf{Y}^*) \approx \sigma(\log p(\pi) + \gamma \sum_{t=\tau}^{T-1} \mathcal{F}(\mathbf{y}_t^*, \mathbf{x}_t, | \pi)). \quad (4)$$

Moreover, for the ‘past’ where $t = [0, \dots, \tau - 1]$, minimizing $\mathcal{F}(\cdot)$ amounts for state estimation of \mathbf{x} given \mathbf{y} . Another consideration is whether the dynamic and observation models are known or unknown. In this work we assume they are given, but AI can also include estimating these models from data.

3.2 Active Inference in Practice

Initial AI work was restricted to discrete domains and evaluated on simple grid-world environments [5, 6]. Later work on continuous state spaces use various black-box approaches such as cross-entropy [25], evolutionary strategies [26], and policy gradient [16] to infer π . A model-based method was achieved by using stochastic VI on expert data [3]. Connections between AI and CaI, performing inference via message passing, have been previously discussed [13, 27]. AI has been applied to real robots for kinematic planning, performing gradient descent on the free energy using the Laplace approximation every timestep [18]. Despite these various approaches, AI has yet to demonstrate the sophisticated control achieved by advanced optimal methods, such as differential dynamic programming [20].

4 Control as Inference

From its origins in probabilistic control design [12], defining a state $\mathbf{z} \in \mathbb{R}^{d_z}$ to describe the desired system trajectory¹ $p(\mathbf{Z})$, optimal control can be expressed as finding the state-action distribution that minimizes the distance for a generative model parameterized by θ , which can be framed as a likelihood objective [17]

$$\min \mathbb{D}_{\text{KL}}[p(\mathbf{Z}) || q_{\theta}(\mathbf{Z})] \equiv \max \mathbb{E}_{\mathbf{Z} \sim p(\cdot)} [\log \int q_{\theta}(\mathbf{Z}, \mathbf{X}, \mathbf{U}) d\mathbf{X} d\mathbf{U}]. \quad (5)$$

When $p(\mathbf{Z})$ simply describes a desired state \mathbf{z}_t^* , so $p(\mathbf{z}_t) = \delta(\mathbf{z}_t - \mathbf{z}_t^*)$, and the latent state-action trajectory is approximated by $q_{\phi}(\mathbf{X}, \mathbf{U})$, the objective (Equation 5) can be expressed as an ELBO where the ‘data’ is \mathbf{Z}^*

$$\max \mathbb{E}_{\mathbf{X}, \mathbf{U} \sim q_{\phi}(\cdot | \mathbf{Z}^*)} [\log q_{\theta}(\mathbf{Z}^* | \mathbf{X}, \mathbf{U})] - \mathbb{D}_{\text{KL}}[q_{\phi}(\mathbf{X}, \mathbf{U} | \mathbf{Z}^*) | q_{\theta}(\mathbf{X}, \mathbf{U})], \quad (6)$$

¹ while \mathbf{z} could be defined from $[\mathbf{x}, \mathbf{u}]^{\top}$, it could also include a transformation, e.g. applying kinematics to joint space-based control for a cartesian space objective.

where ϕ captures the latent state parameterization and θ defines the remaining terms, i.e. the priors on the system parameters and latent states. This objective can be optimized using EM, estimating the latent state-action trajectory ϕ in the E step and optimizing the remaining unknowns θ in the M step. By exploiting the temporal structure, $q_\phi(\mathbf{X}, \mathbf{U} \mid \mathbf{Z}^*)$ can be inferred efficiently in the E step by factorizing the joint distribution (Equation 7) and applying Bayes rule recursively

$$q_\phi(\mathbf{Z}^*, \mathbf{X}, \mathbf{U}) = q_\phi(\mathbf{x}_0) \prod_{t=0}^{T-1} q_\phi(\mathbf{x}_{t+1} \mid \mathbf{x}_t, \mathbf{u}_t) \prod_{t=0}^T q_\phi(\mathbf{z}_t^* \mid \mathbf{x}_t, \mathbf{u}_t) q_\phi(\mathbf{u}_t \mid \mathbf{x}_t), \quad (7)$$

$$q_\phi(\mathbf{x}_t, \mathbf{u}_t \mid \mathbf{z}_{0:t}^*) \propto q_\phi(\mathbf{z}_t^* \mid \mathbf{x}_t, \mathbf{u}_t) q_\phi(\mathbf{x}_t, \mathbf{u}_t \mid \mathbf{z}_{0:t-1}^*), \quad (8)$$

$$q_\phi(\mathbf{x}_t, \mathbf{u}_t \mid \mathbf{z}_{0:T}^*) \propto q_\phi(\mathbf{x}_t, \mathbf{u}_t \mid \mathbf{x}_{t+1}) q_\phi(\mathbf{x}_{t+1} \mid \mathbf{z}_{0:T}^*). \quad (9)$$

Equations 8, 9 are commonly known as Bayesian filtering and smoothing [19]. The key distinction of this framework from state estimation is the handling of \mathbf{u} during the forward pass, as $q_\phi(\mathbf{x}_t, \mathbf{u}_t) = q_\phi(\mathbf{u}_t \mid \mathbf{x}_t) q_\phi(\mathbf{x}_t)$, control is incorporated into the inference. We can demonstrate this in closed-form with linear Gaussian inference and linear quadratic optimal control.

4.1 Linear Gaussian Inference & Linear Quadratic Control

While the formulation above is intentionally abstract, it can be grounded clearly by unifying linear Gaussian dynamical system inference (LGDS, i.e. Kalman filtering and smoothing) and linear quadratic Gaussian (LQG) optimal control [22]. While both cases have linear dynamical systems, here LQG is fully-observed² and has a quadratic control cost, while the LGDS is partially observed and has a quadratic log-likelihood due to the Gaussian additive uncertainties. These two domains can be unified by viewing the quadratic control cost function as an Gaussian observation likelihood. For example, given $\mathbf{z}_t = \mathbf{x}_t + \boldsymbol{\xi}$, $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ and $\mathbf{z}_t^* = \mathbf{0} \forall t$,

$$\log q_\theta(\mathbf{z}_t^* \mid \mathbf{x}_t, \mathbf{u}_t) = -\frac{1}{2}(d_z \log 2\pi + \log |\boldsymbol{\Sigma}| + \mathbf{x}_t^\top \boldsymbol{\Sigma}^{-1} \mathbf{x}_t) = \alpha \mathbf{x}_t^\top \mathbf{Q} \mathbf{x}_t + \beta \quad (10)$$

where (α, β) represents the affine transformation mapping the quadratic control cost $\mathbf{x}^\top \mathbf{Q} \mathbf{x}$ to the Gaussian likelihood. As convex objectives are invariant to affine transforms, this mapping preserves the control problem while translating it into an inference one. The key unknown here is α , which incorporates \mathbf{Q} into the additive uncertainty $\boldsymbol{\xi}$, $\boldsymbol{\Sigma} = \alpha \mathbf{Q}^{-1}$. Moreover, inference is performed by using message passing [15] in the E step to estimate \mathbf{X} and \mathbf{U} , while α is optimized in the M step. This view scales naturally to not just the typical LQG cost $\mathbf{x}^\top \mathbf{Q} \mathbf{x} + \mathbf{u}^\top \mathbf{R} \mathbf{u}$, but also nonlinear mappings to \mathbf{z} by using approximate inference. While the classic LQG result includes the backward Riccati equations

² Confusingly, LQG can refer to both Gaussian disturbance and/or observation noise. While all varieties share the same optimal solution as LQR, the observation noise case results in a partially observed system and therefore requires state estimation. i2c is motivated by the LQR solution and therefore does not consider observation noise, but it would be straightforward to integrate.

and an optimal linear control law, the inference setting derives direct parallels to the backward pass during smoothing [22] and the linear conditional distribution of the Gaussian, $q_{\theta}(\mathbf{u}_t | \mathbf{x}_t) = \mathcal{N}(\mathbf{K}_t \mathbf{x}_t + \mathbf{k}_t, \Sigma_{k_t})$ [10] respectively. As the conditional distribution is linear, updating the prior joint density $p(\mathbf{x}_t, \mathbf{u}_t)$ in the forward pass with updated state estimate \mathbf{x}'_t corresponds to linear feedback control w.r.t. the prior

$$p(\mathbf{u}'_t) = \int p(\mathbf{u}_t | \mathbf{x}_t = \mathbf{x}'_t) p(\mathbf{x}'_t) d\mathbf{x}'_t, \quad (11)$$

$$\boldsymbol{\mu}_{\mathbf{u}'_t} = \boldsymbol{\mu}_{\mathbf{u}_t} + \mathbf{K}_t (\boldsymbol{\mu}_{\mathbf{x}_t} - \boldsymbol{\mu}_{\mathbf{x}'_t}), \quad (12)$$

$$\Sigma_{\mathbf{u}\mathbf{u}'_t} = \Sigma_{\mathbf{u}\mathbf{u}_t} - \Sigma_{\mathbf{u}\mathbf{x}_t} \Sigma_{\mathbf{x}\mathbf{x}_t}^{-1} \Sigma_{\mathbf{x}\mathbf{u}_t}^{\top} + \mathbf{K}_t \Sigma_{\mathbf{x}\mathbf{x}'_t} \mathbf{K}_t^{\top}, \quad (13)$$

$$\mathbf{K}_t = \Sigma_{\mathbf{u}\mathbf{x}_t} \Sigma_{\mathbf{x}\mathbf{x}_t}^{-1}. \quad (14)$$

From Equation 14, it is evident that the strength of the feedback control depends on both the certainty in the state and the correlation between the optimal state and action.

The general EM algorithm for obtaining $q_{\theta}(\mathbf{x}, \mathbf{u})$ from $p(\mathbf{Z})$ is referred to as input inference for control (I2C) [28] due to its equivalence with input estimation. Note that for linear Gaussian EM, the ELBO is tight as the variational distribution is the exact posterior. For nonlinear filtering and smoothing, mature approximate inference methods such as Taylor approximations, quadrature and sequential Monte Carlo may be used for efficient and accurate computation [19].

Another aspect to draw attention to is the inclusion of \mathbf{z} compared to alternative CaI formulations, which frame optimality as the probability for some discrete variable \mathbf{o} , $p(\mathbf{o}=1 | \mathbf{x}, \mathbf{u})$ [14]. Previous discussion on CaI vs AI have framed this discrete variable as an important distinction. However, it is merely a generalization to allow for a general cost function $C(\cdot)$ to be framed as a log-likelihood, i.e. $p(\mathbf{o}=1 | \mathbf{x}, \mathbf{u}) \propto \exp(-\alpha C(\mathbf{x}, \mathbf{u}))$. For the typical state-action cost functions that are a distance metric in some transformed space, the key consideration is the choice of observation space \mathbf{z} and corresponding exponential density.

5 The Unifying View: Control of the Observations

A key distinction to the AI and CaI formulations described above is that, while AI combines state estimation and control with a unified objective, CaI focuses on trajectory optimization. However, this need not be the case. In a similar fashion to the partially-observed case of LQG, CaI also naturally incorporates observations [23]. As Section 4 describes I2C through a general Bayesian dynamical system, the formulation can be readily adapted to include inference using past measurements. Moreover, as I2C frames the control objective as an observation likelihood, when \mathbf{z} and \mathbf{y} are the same transform of \mathbf{x} and \mathbf{u} , the objective can also be unified and directly compared to active inference. For ‘measurements’

$\mathbf{Y}^* = \{\mathbf{y}_0^*, \dots, \mathbf{y}_{\tau-1}^*, \mathbf{z}_\tau^*, \dots, \mathbf{z}_{T-1}^*\}$, following Equation 5 using the $\mathcal{F}(\cdot)$ notation

$$\min \mathbb{D}_{\text{KL}}[p(\mathbf{Y})||q_\theta(\mathbf{Y})] = \min \underbrace{\sum_{t=0}^{\tau-1} \mathcal{F}_\psi(\mathbf{y}_t^*, \mathbf{x}_t, \mathbf{u}_t)}_{\text{state estimation}} + \underbrace{\sum_{t=\tau}^{T-1} \mathcal{F}_\psi(\mathbf{z}_t^*, \mathbf{x}_t, \mathbf{u}_t)}_{\text{optimal control}}, \quad (15)$$

where $\psi = \{\theta, \phi\}$. Here, $p(\mathbf{y}_t) = \delta(\mathbf{y}_t - \mathbf{y}_t^*)$ now also describes the empirical density of past measurements $\mathbf{y}_{<\tau}^*$. The crucial detail for this representation is that the observation model $q_\theta(\mathbf{y}_t | \mathbf{x}_t, \mathbf{u}_t, t)$ is now time dependent, switching from estimation to control at $t = \tau$. For the Gaussian example in Section 4.1, $\Sigma_{<\tau}$ is the measurement noise and $\Sigma_{\geq\tau}^{-1} = \alpha \mathbf{Q}$. A benefit of this view is that the computation of active inference can now be easily compared to the classic results of Kalman filtering and LQG (Fig. 1), and also scaled to nonlinear tasks through approximate inference. Moreover, obtaining the policy $\pi(\cdot)$ using the joint distribution $q_\theta(\mathbf{x}_t, \mathbf{u}_t)$ is arguably a more informed approach compared to direct policy search on an arbitrary policy class.

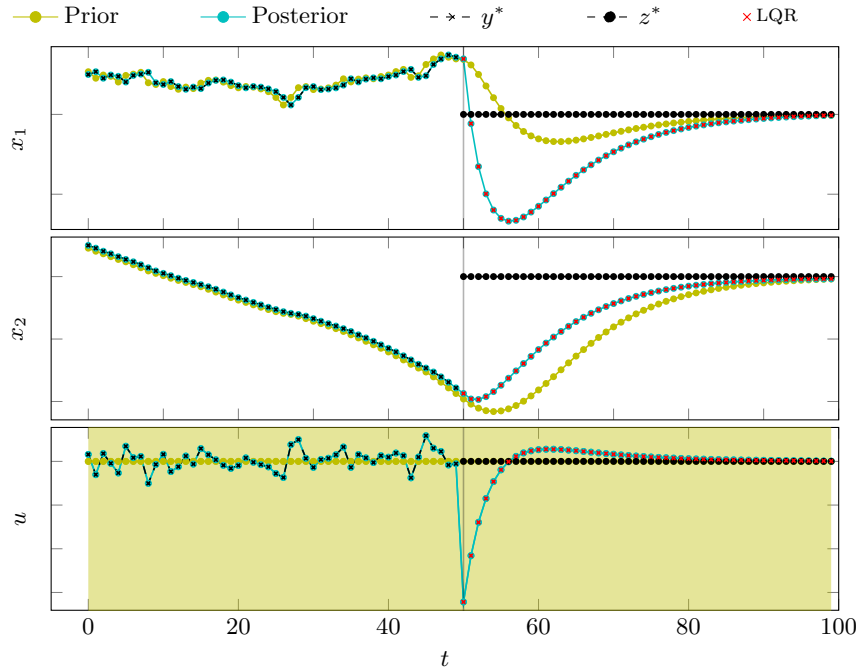


Fig. 1. Linear Gaussian i2C performing state estimation and control following Section 5, with state $\mathbf{x}=[x_1, x_2]^\top$, action u and $[\mathbf{x}, u]^\top$ as the observation space. With $\tau = 50$, for $t < \tau$ i2C performs state estimation under random controls. For $t \geq \tau$, i2C switches to optimal control. This example is in the low noise setting, with a large prior on u , to illustrate that i2C returns the LQR solution for the same initial state and planning horizon.

6 Conclusion

We have derived an equivalent formulation to active inference by considering partially-observed, inference-based optimal control, which has a principled derivation and is well-suited for approximate inference. While we have delineated state estimation as operating on past measurement and control as planning future actions (Equation 15), both AI and I2C demonstrate the duality between estimation and control due to the mathematical similarity when both are treated probabilistically. We hope the inclusion of the CaI literature enables a greater theoretical understanding of AI and more effective implementations through approximate inference.

References

1. Attias, H.: Planning by probabilistic inference. In: International Workshop on Artificial Intelligence and Statistics (2003)
2. Biehl, M., Guckelsberger, C., Salge, C., Smith, S.C., Polani, D.: Expanding the active inference landscape: More intrinsic motivations in the perception-action loop. *Frontiers in Neurorobotics* (2018)
3. Catal, O., Nauta, J., Verbelen, T., Simoens, P., Dhoedt, B.: Bayesian policy selection using active inference. In: ICLR Workshop on Structure & Priors in Reinforcement Learning (2019)
4. Friston, K.: The free-energy principle: a unified brain theory? *Nature reviews neuroscience* (2010)
5. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., Pezzulo, G.: Active inference: a process theory. *Neural computation* (2017)
6. Friston, K.J., Daunizeau, J., Kiebel, S.J.: Reinforcement learning or active inference? *PloS one* (2009)
7. Gershman, S.J.: What does the free energy principle tell us about the brain? *Neurons, Behavior, Data analysis, and Theory* (2019)
8. Guzmán, N.: twitter.com/NoahGuzman14/status/1259953086241492992, Accessed: 2020-06-17
9. Herreros, I., Verschure, P.F.: About the goal of a goals' goal theory. *Cognitive Neuroscience* (2015)
10. Hoffmann, C., Rostalski, P.: Linear optimal control on factor graphs - a message passing perspective -. *International Federation of Automatic Control* (2017)
11. Kappen, H.J.: Path integrals and symmetry breaking for optimal control theory. *Journal of Statistical Mechanics: Theory and Experiment* (2005)
12. Kárný, M.: Towards fully probabilistic control design. *Automatica* (1996)
13. van de Laar, T., Özçelikkale, A., Wymeersch, H.: Application of the free energy principle to estimation and control. *arXiv preprint arXiv:1910.09823* (2019)
14. Levine, S.: Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909* (2018)
15. Loeliger, H.A., Dauwels, J., Hu, J., Korl, S., Ping, L., Kschischang, F.R.: The factor graph approach to model-based signal processing. *Proc. of the IEEE* (2007)
16. Millidge, B.: Deep active inference as variational policy gradients. *arXiv preprint arXiv:1907.03876* (2019)
17. Murphy, K.P.: *Machine learning: a probabilistic perspective*. MIT press (2012)

18. Oliver, G., Lanillos, P., Cheng, G.: Active inference body perception and action for humanoid robots. arXiv preprint arXiv:1906.03022 (2019)
19. Srkk, S.: Bayesian Filtering and Smoothing. Cambridge University Press (2013)
20. Tassa, Y., Erez, T., Todorov, E.: Synthesis and stabilization of complex behaviors through online trajectory optimization. In: International Conference on Intelligent Robots and Systems. IEEE (2012)
21. Todorov, E.: Linearly-solvable markov decision problems. In: Advances in neural information processing systems (2007)
22. Toussaint, M.: Robot trajectory optimization using approximate inference. In: International conference on machine learning (2009)
23. Toussaint, M., Charlin, L., Poupart, P.: Hierarchical pomdp controller optimization by likelihood maximization. In: Uncertainty in Artificial Intelligence (2008)
24. Toussaint, M., Storkey, A.: Probabilistic inference for solving discrete and continuous state Markov Decision Processes. In: International Conference on Machine Learning (2006)
25. Tschantz, A., Millidge, B., Seth, A.K., Buckley, C.L.: Reinforcement learning through active inference. arXiv preprint arXiv:2002.12636 (2020)
26. Ueltzhöffer, K.: Deep active inference. Biological Cybernetics (2018)
27. de Vries, B., Friston, K.J.: A factor graph description of deep temporal active inference. Frontiers in Computational Neuroscience (2017)
28. Watson, J., Abdulsamad, H., Peters, J.: Stochastic optimal control as approximate input inference. In: Conference on Robot Learning (2019)